Operation

Apart from far more sophisticated error-checking techniques, the innards of a CD-ROM drive are pretty much the same as those used in CD audio players. Data is stored in the same way on all CDs. The information is stored in sequential 2KB sectors that form a single spiral track that starts at the centre of the disc and wraps around many times until it reaches the outer edge of the disc.

A player reads information from the CD’s spiral track of pits and lands, starting from the centre of the disc and moving to the outer edge. It does this by firing an infrared laser - 780 nano-millimetres wide and generated by a small gallium arsenate semiconductor - through the clear optical grade polycarbonate plastic layer and onto the metallic sheet. Although itís of very low power, itís strong enough to damage the eye if shined directly into it. As the disc rotates at between 200 and 500rpm, the light bounces off the pits and the frequency of the light changes.

The areas around the pits, called lands, also play a part in the process. The reflected light passes through a prism and onto a photosensor, the output of which is proportional to the amount of light it receives. Light reflected from a pit is 180 degrees out of phase with the light from the lands, and the differences in intensity are measured by the photo-electric cells and converted into electrical pulses. The result is that the series of pits and lands of varying lengths stamped into the surface of the disc are interpreted as a series of corresponding 0s and 1s from which the data - or, via a digital-to analogue converter (DAC), the audio - stored on the disc is recreated. And since just a weak bandwidth laser is the only thing to touch the surface of the CD directly, thereís none of the wear and tear of traditional analogue media to contend with.

Things would be relatively simple if CD-ROM discs were perfectly flat and could be spun with no horizontal deviation. In fact, a considerable amount of extra electronics wizardry is needed to ensure that the laser stays in focus on the disc surface and that it follows the track itís reading.

There are various methods for maintaining radial tracking, the most common being the three-beam approach. The laser beam isnít shone directly onto the disc surface but is emitted from a semiconductor laser unit and passed through a diffraction grating to produce two extra light sources, one each side of the main beam. A collimator lens takes the three beams and makes them parallel, after which theyíre passed through a prism called a polarised beam splitter. The beam splitterís job is to allow the outbound beams to pass through while reflecting the returning beams through 90 degrees down to the photodiode that interprets the signal.

The two side beams are measured for intensity, which remains equal as long as they stay on either side of the track. Any sideways movement of the disc will result in an imbalance and a servo motor will reposition the objective lens. Vertical movement is countered by splitting the receptor photodiode into four quadrants and placing it halfway between the horizontal and vertical focal points of the beam. Any deviation of the disc will cause the spot to become elliptical, with a corresponding current imbalance between each opposing pair of quadrants. The objective lens is then moved up or down to ensure the spot remains circular.

CD technology has built-in error correction systems which are able to suppress most of the error that arise from physical particles on the surface of a disc. Every CD-ROM drive and CD player in the world uses Cross Interleaved Reed Solomon Code (CIRC) detection and the CD-ROM standard provides a second level of correction via the Layered Error Correction Code algorithm. With CIRC, an encoder adds two dimensional parity information, to correct errors, and also interleaves the data on the disc to protect from burst errors. It is capable of correcting error bursts up to 3,500 bits (2.4 mm in length) and compensates for error bursts up to 12,000 bits (8.5 mm) such as caused by minor scratches.

Digital audio

On vinyl and audio cassettes, the audio waveform is recorded as an analogue signal. Therefore any imperfections will be heard as noise (hiss) or other defects. To reduce these defects, CDs use digital techniques, storing ‘samples’ as numbers. The process of converting analogue to digital is known as digitising or sampling. The analogue waveform is chopped into a number of slices per second. At each slice, the amplitude is measured and rounded to the nearest available value. Clearly the more chops per second (sampling rate) and the finer the values assignable to the amplitude (dynamic range), the better the representation of the original.

CD digital employs a sampling rate of 44.1kHz and a 16-bit dynamic range. That is, 44,100 chops every second, each one describing the waveform amplitude at that moment in time with a 16-bit number; 16-bit itself offering 65,536 steps from which to choose. This sampling rate provides a frequency response adequate for sounds up to 20kHz in pitch - although some audiophiles argue that it isn't sufficient to capture essential psychoacoustic effects that occur beyond the range of human hearing. Conventional wisdom holds that 16-bits provides ample dynamic range for loud and soft musical passages. The audio is recorded on two tracks, for stereo sound.

A 44.1kHz rate means there are 44,100 chops every second, each one describing the waveform amplitude at that moment in time with a 16-bit number; 16-bit itself offering 65,536 steps from which to choose. With samples occupying two bytes on each of two channels - each sample yields a data transfer rate of just over 176KBps. A single-speed CD-ROM transfers data at the same rate, but a portion of the data stream is taken by error-correcting information - reducing the effective transfer rate to 150KBps. A CD can hold up to 74 minutes of encoded stereo audio data - which, when the ECC overhead is taken account of, equates to the standard CD capacity of 680MB.

The table below summarises the relevant parameters:

Sample rate	44.1kHz
Channels	2 (stereo)
Bits per sample, per channel	16
Levels per sample	65,536
Total data rate	176KBps
Effective data rate	150KBps
Effective CD capacity	680MB
CD playing time	74 minutes

CLV

The first generation of ësingle-speedí CD-ROM drives were based on the design of audio CD drives, employing constant linear velocity (CLV) technology to spin a disc at the same speed as an audio CD which, with error correction, meant 150KBps.

Since there are more sectors on the outside edge of the CD than in the centre CLV uses a servo motor to slow the spin speed of the disc towards the outer tracks in order to maintain a constant data transfer rate over the laser read head. The internal memory buffer of the drive controls this by using a quartz crystal oscillator to clock the data out of the buffer at a specific rate and maintain it at a 50% capacity while data is being read into it. If the data is being retrieved too fast, the 50% capacity threshold will be breached, so a command is sent to slow the drive motor down.

While audio discs have to be read at single-speed, there is no limit for the CD-ROM. Indeed, the faster the data is read off, the better. As CD-ROM technology progressed, speeds increased at regular and seemingly ever-shorter intervals until, by early 1998, the fastest drives were capable of 32-speed, a data transfer rate of 4.8MBps.

A quad-speed drive employing CLV technology, for example, span the disc at around 2,120rpm when reading the inner tracks and 800rpm when reading the outer tracks. Variable spin is also necessary to cope with audio data, which is always read at single-speed (150KBps), whatever the transfer rate for computer data. The key issues with spin speed, then, are the quality of the spindle motor that rotates the disc and of the software that controls its operation at varying rotational velocities, as well as the ability of the read head positioning system to move quickly and accurately into position to begin accessing data. Unless the driveís engineering and electronics are up to the job, merely increasing the basic spin speed wonít deliver the performance boost which might otherwise be expected.

A related factor is the level of utilisation of the host PCís CPU: as the spin speed, and therefore the data transfer rate, increases, so does the amount of the CPU attention required to process the data streaming off the CD-ROM. This means that if other tasks are accessing the processor at the same time, then thereíll be less processing power available for the CD-ROM drive, and transfer rates will fall. A well-designed CD-ROM drive system should seek to minimise the CPU utilisation for a given spin speed and data transfer rate. If it doesnít a fast CD-ROM drive might have inferior overall performance than a slower-speed model.

The size of data buffer is another widely-quoted statistic; a big on-drive buffer - 1MB as opposed to 128K, say - is certainly desirable for boosting burst data throughput, but unless this is controlled by high-quality firmware, the resulting small performance increase will hardly justify the cost of the extra buffer memory.

CAV

CLV remained the dominant CD-ROM technology until Pioneer, the first with quad-speed when most only had single, skipped six- and eight-speed entirely in favour of the new ten-speed speed technology employed by its DR-U10X drive, launched in mid-1996. The clever part was that the drive operated not only in conventional CLV (constant linear velocity) mode, but also CAV (constant angular velocity) mode, transferring data at a variable rate while the drive spins at a constant rate, like a hard disk.

Access time has a big influence on overall performance. As the speed of a CLV drive increases, access times often suffer as it becomes harder to perform the abrupt changes in spindle velocity needed to maintain a constant high data transfer rate, due to the mass inertia of the disc itself. CAV maintains a steady spin speed resulting in increased data transfer rates and reduced seek times as the head moves towards its outside edge. While early CLV drives had average access times greater than 500 ms, modern CAV drives typically have average access times less than 100 ms.

Pioneerís revolutionary design allowed operation exclusively in either CLV or CAV mode, or in a mixed mode. In the latter, CAV was used for reading close to the centre of the disc while the drive switched to CLV mode for reads closer to the outer edge. The Pioneer drive signalled the end of the line for CLV-only designs and the so-called Partial CAV drives assumed the mantle of the leading CD-ROM technology.

This remained the case until the advent of a new generation of DSP (digital signal processing) chips capable of handling faster than 16-speed data rates and, in the autumn of 1997 Hitachi launched the first CD-ROM drive to use Full CAV technology. This overcame many of the problems inherent in Partial CAV designs, obviating the need to monitor the head location and step the motor speed up and down to maintain a steady DTR and making access times more consistent because there is no waiting for the spindle speed to settle between transitions.

Most of the 24-speed full CAV CD-ROMs on the market by the end of 1997 spun discs at a constant 5,000rpm, achieving a DTR of 1.8MBps at the centre increasing to 3.6MBps at the outer edge. By the summer of 1999 outer track DTRs had been increased to 48-speed, or 7.2MBps, as a result of increasing spindle speeds to an incredible 12,000 rpm - comparable to the rotation speed of many high-performance hard disk drives.

However, one of the main problems with spinning CD-ROMs at such high speeds is excessive noise and vibration, often including loud hissing noise caused by air being forced out of the drive casing by the spinning CD-ROM. Because the CD-ROM is clamped at its centre, the most severe vibration occurs at the outer edges of the disc - at exactly the point where the decoding circuits have to handle the highest signal rate. Also, since few CD-ROMs actually have data stored on their outer edges these high spin-rate drives rarely achieve their theoretical maximum DTR in real-life.

Applications

Then there is the question of which applications actually benefit from a faster CD-ROM drive. Most multimedia titles are optimised for double or, at best, quad-speed drives. If video is recorded to play back in real time at a 300KBps sustained transfer rate, anything faster than double-speed is unnecessary. In some cases, a faster drive may be able to read off the information quickly into a buffer cache, from where it is subsequently played, freeing up the drive for further work. This is rare, however.

Pulling off large images from a PhotoCD would be a perfect application for a faster CD-ROM drive, but decompressing these images as theyíre read off the disc results in a performance ceiling of quad-speed. In fact, just about the only application which truly needs a fast data transfer rate is copying sequential data onto a hard disc; in other words, installing software.

Fast CD-ROM drives are only fast for sustained data transfer, not random access. An ideal application for high sustained data transfer is high-quality digital video, recorded at a suitably high rate. MPEG-2 video, as implemented on Digital Versatile Discs (DVDs), requires a sustained data transfer of around 580KBps, compared to MPEG-1ís 170KBps found on existing White Book VideoCDs. However, a standard 650MB CD-ROM disc would last less than 20 minutes at those high rates, so high-quality video will only be practical on DVD discs, which have a much higher capacity.

Interfaces

CD-ROM drives have three main connections on the back: power, audio out to a sound card and a data interface.

These days it is increasingly common to find CD-ROM drives with an IDE data interface, which in theory can be connected to the IDE controller found on virtually every PC. The original IDE hard drive was designed for the AT Bus, and the older IDE interface allowed two hard drives to be attached, one as master and the other as slave. Later, the ATAPI specification allowed one of these to be an IDE CD-ROM drive. EIDE takes this a stage further by adding a second IDE channel for two more devices, and allows a mixture of hard discs, CD-ROMs and tape drives to be used.

An operation on one of these devices must be completed before any others can be accessed. Putting a CD-ROM on the same channel as a hard drive will affect performance as its much slower and will block access to the disc. In systems with two IDE hard discs, the CD-ROM drive should be isolated on the secondary IDE channel while the discs are fitted as master and slave on the primary channel. The discs have to contend with each other but at least the CD-ROM drive wonít get in the way. Other drawbacks of EIDE are that the number of devices which can be attached is limited to four and devices must be mounted internally, so expansion can be limited by the size of the PC.

The SCSI-2 standard allows a maximum of 14 devices can be attached to one host adaptor card and these may be a mixture of internal and external devices. SCSI allows all devices on the bus to be active simultaneously, although only one can be transmitting data. Physically locating data on devices is more time-consuming, so while one device is using the SCSI bus, any other devices can be positioning their heads for read or write operations. The latest Fast Wide SCSI specification supports a maximum data transfer rate of 20MBps compared to EIDEís 13MBps and, as SCSI devices have more built-in intelligence they are far less CPU intensive than IDE devices.

SCSI scores over IDE again with its more frugal use of the PCís resources - namely IRQs. Due to the number of extra cards and devices, todayís multimedia, Internet-ready, networked PCs make heavy demands on IRQs, leaving little or no room for further upgrades. The primary EIDE interface is usually allocated IRQ14, while the secondary gets IRQ15, so four devices can be added for the cost of two interrupts. SCSI is far less resource hungry, as no matter how many devices are attached to the bus, only a single IRQ is required for the host adaptor.

In summary, SCSI has greater expansion potential and better performance, but comes at a much higher cost than IDE. The current preponderance of EIDE internal drives would therefore seem to be more a matter of convenience and cost rather than technical superiority and SCSI remains the interface of choice for external CD-ROM drives.

DMA vs PIO mode

Traditionally, CD-ROM drives have used Programmable Input Output (PIO) rather than Direct Memory Access (DMA) for data transfer. This was favoured for the earlier designs because hardware implementation is simpler and adequate for devices that require low transfer rates. The drawback is that the CPU must mediate the transfer of data, often byte by byte. As the data rate of CD-ROM drives has risen, so has the load on the CPU, to the point where 24-speed and 32-speed drives can completely saturate CPU utilisation in PIO mode. The severity of CPU loading depends on a number of factors, such as the exact PIO mode used, the PC's IDE/PCI bridge design, the CD-ROM buffer size and design and the CD-ROM device driver.

DMA data transfer is always more efficient and requires only a few per cent of CPU time. It uses hardware to control data transfer directly to system memory, and only require initial memory allocation and minimal handshaking from the CPU. A further advantage is that performance is device rather than system dependent. DMA-capable devices should give consistent performance, regardless of the system they're attached to.

DMA has been a standard feature of most SCSI systems for some time, but has only recently been common for IDE devices and interfaces. Windows has been a major drawback to DMA implementation in many desktop PCs, and it's only the recent Windows 95 OSR2 and OSR2.1 OEM service releases that have included DMA drivers for the Intel PIIX range of IDE/PCI bridge chips but DMA is often not enabled in systems equipped with these operating system versions, either due to ignorance or from concerns about overall system stability.

TrueX technology

Based on the aim of enabling users to run application direct from CD without first cacheing to the hard disk drive, Zen Research's TrueX technology takes a different route to increasing the performance of CD-ROM drives - addressing data transfer rates and access times rather than simply seeking to spin discs faster. Unlike the conventional CD-ROM, which uses a highly concentrated laser beam to read the digital signal that's encoded as tracks of small pits on the surface of a disc, the Zen approach uses a custom ASIC to illuminate multiple tracks, detect them simultaneously and read them in parallel. The ASIC integrates analogue interface elements like digital phase-locked loop (DPLL), DSP, servo-motor controller, parallel to serial data converter, data decoding and error detection/correction and ATAPI interface. Provisions have been made to allow the connection of an external SCSI or IEEE 1394 interface chip if required.

A diffracted laser beam used in conjunction with a multiple beam detector array illuminates and detects multiple tracks. A conventional laser diode is sent through a diffraction grating that splits the beam into seven discrete beams, spaced evenly to illuminate seven tracks. The seven beams pass through the beam splitting mirror to the objective lens and onto the surface of the disc. Focus and tracking are attained with the central beam. Three beams on either side of the centre are readable by a detector array as long as the centre is on track and in focus. The reflected beams return via the same path and are directed to the detector array by the beam splitter mirror. Seven detectors are contained in the multi-beam detector to align with the reflected tracks. Conventional detectors are also provided for focus and tracking.

Although the mechanical elements of the CD-ROM drive are slightly changed - disc rotation and read head motion remain the same - the format of the disc media remains CD or DVD standard and the design uses a conventional approach to tracking and seeks. Although TrueX can be applied to both CLV and CAV disc rotation systems, Zen has focused on CLV in order to deliver constant transfer rates across the whole disc. In either case, higher data rates are achieved at lower, more disc-tolerant rotation speeds, reducing vibration and promoting smoother operation and greater reliability.

Kenwood Technologies shipped its first TrueX CD-ROM drive - a 40-speed device - in August 1998 and followed this up six months later by unveiling a 52-speed device. A typical CD is less than half full, and the majority of all CDs contain no data on their outer tracks. Depending upon the operating environment and quality of media, the Kenwood 52X TrueX CD-ROM drive delivers a typical performance ranging from 6750 - 7800KBps (45X to 52X) across the entire disc. By comparison, a conventional 48X 'max' CD-ROM drive performs at 19X on its innermost tracks, but achieves 48X performance only on the disc's outermost tracks (if the disc is full) - and at rotational speeds more than double that of the Kenwood device!